HPS Computing Needs Planning

SCA planning meeting

15 February 2012

- H. Neal

Figure 2.2.1.2. Layout of the HPS test run chicane and detector package.

What HPS is about

Heavy photons in this mass/coupling range are expected on very general theoretical grounds, and also motivated by recent astrophysical evidence suggesting they might mediate dark matter annihilations and/or dark matter interactions with ordinary matter.

The physics which motivates the HPS Test Run is exactly that which motivates the full HPS, and is discussed in detail in the HPS proposal [1]. Briefly, HPS is searching for new heavy vector boson(s), aka "heavy photons" or "dark photons" or "hidden sector photons", in the mass range of 20 MeV/c^2 to 1000 MeV/c^2. Heavy photons mix with the Standard Model photon through kinetic mixing, which induces their weak coupling to electrons, εe, where $\varepsilon \sim 10^{-3}$. Heavy photons in this mass/coupling range are expected on very general theoretical grounds, and also motivated by recent astrophysical evidence suggesting they might mediate dark matter annihilations and/or dark matter interactions with ordinary matter. Since they couple to electrons, heavy photons are radiated in electron scattering and can subsequently decay into narrow e^+e^- resonances which can be observed above the copious QED trident background. For suitably small couplings, heavy photons travel detectable distances before decaying, providing a second signature. The HPS experiment exploits both of these signatures to search for heavy photons over a wide range of couplings, $\varepsilon^2 > 10^{-10}$, and masses.

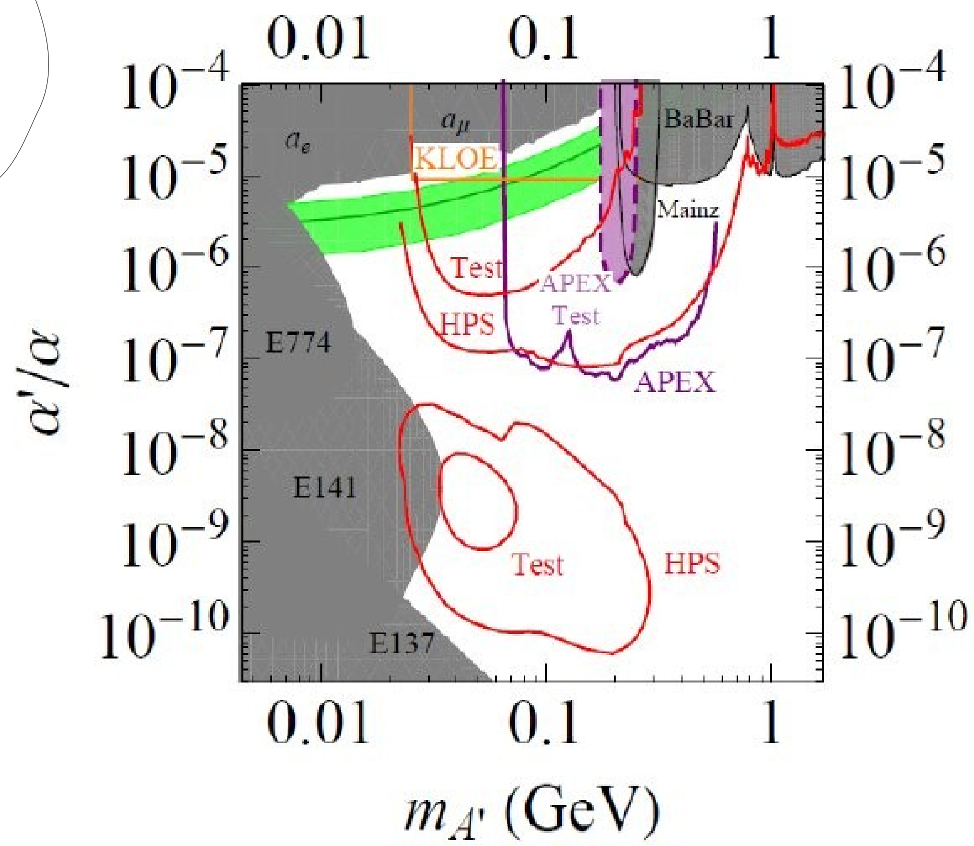

HPS Test Run

The highest priority goal of the HPS Test run is to validate the fundamental assumptions behind the design of the full HPS experiment and thereby prepare the foundation for full HPS approval.

Scheduled for two weeks of running in March/April 2012

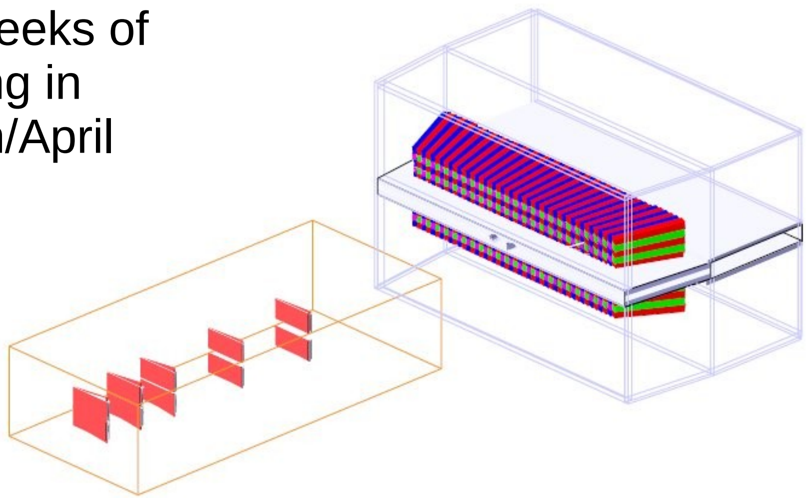

Figure 3.1.1.2. Rendering of the detector simulation used for trigger studies. The orange box represents the container for the silicon trackers (outlined only), the red rectangles represent the silicon tracking layers, and the larger blue-gray outlined rectangles represent the box for the calorimeter. The final object shows the calorimeter, with the crystals colored in alternate colors for clear visibility.

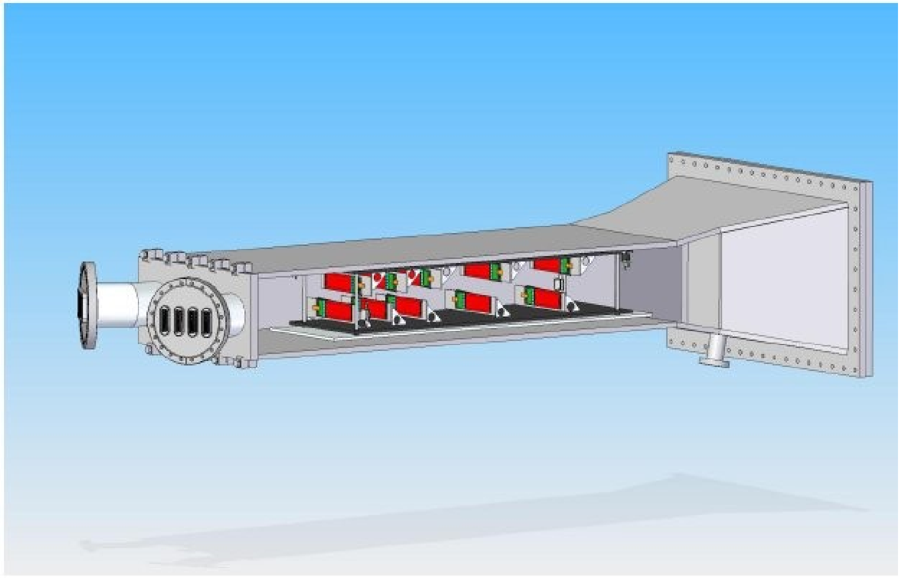

Figure 2.3.6.1: A rendering of the tracker installed inside the existing vacuum chamber from the Hall B TPE analyzing magnet. The new upstream flanges that accommodate connections through the wall of the vacuum chamber are shown at the left.

3.1 **Heavy Photon Signal**

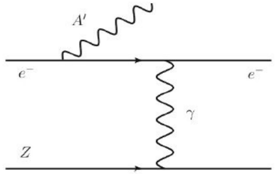

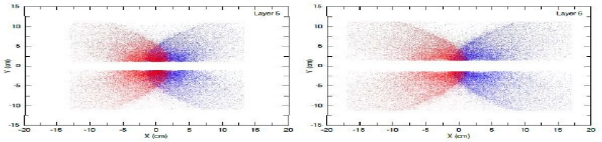

Figure 4.3.4.2: Hits in the silicon tracker in all six layers produced by A' decays with $m_{A'}$ = 300 MeV/c^2 and E_{beam} = 5.5 GeV. All tracks that hit the first five layers are shown, resulting in a five-layer A' acceptance of 44% in the 1T magnetic field. Over 90% of tracks within the five-layer acceptance also produce a hit in Layer 6.

The SVT is the primary data producer

SVT Data Rate and Volume
The amount of data for a ~40 MHz clock and a 50 KHz rate (APV25) is estimated as:
– Structure
HPS Test Run: A proposal to Search for Massive Photons at Jefferson Laboratory
40
• Each chip 12-bit header (3 start, 8-bit address, 1 error) plus 10 bits/channel if 10-bit ADC is used
• For each ASIC need additional 9 bit chip address
– Number of bits for each hit
• 9-bit chip address (16-bit)
• 12-bit header (16-bit)
• 6 times 10-bit ADC value (72-bit) (assume 10-bit ADC for now since need bare minimum 8 bits)
• 13 bytes (104 bits) (81 bit minimum)
– Data volume and rate
• 217 ASICs x 2 bytes chip address: 434 bytes
• 217 hits x 2 bytes header: 434 bytes
• 217 hits x 6 samples x 1.5 bytes: 1953 bytes

The resulting data rate is 2821 bytes x 50 KHz = 141.1 Mbytes/sec.

Maximum trigger rate is 50 KHz

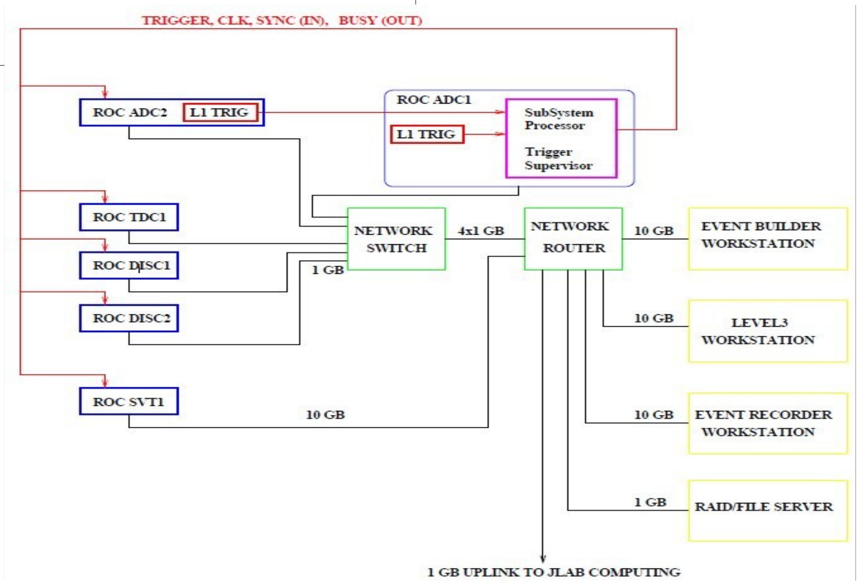

Production needs:

The raw data event sizes per subdetector are provided by the test run proposal at:
https://confluence.slac.stanford.edu/download/attachments/86676777/HPSTestRunProposal-February18.pdf?version=1&

	Event Size (bytes)
ECAL	504
SVT	2821
Total	3325

The trigger rate expected for the first week is 10KHz during commissioning and then upto 50KHz. 30KHz is used as the averaget trigger rate:

	frequency	pass thru rate	Hz
Trigger run1	30000		30000
L3 run1		1	30000
			Bytes/sec
Data Output rate run1			9.98E+07

The net number events for the requested two week run time is:

	runtime(days) per run	uptime fraction	Events
total events	14	50.00%	3.63E+10

Note that I have included a 50% inefficiency since it is very unlikely that we would achieve the maximum data rate early in this two week test run.

Thus the total raw data volume expected is:

Total Data Volume	
60.33	Tbytes

The estimated reconstructed data volume is:

	passing fraction				
Filter	100.00%			3.63E+10	events per pass
		X raw	Size	Net Size	
Recon Data		2	60.33	120.66	Tbytes
MC Data		0.1	90.49	9.05	Tbytes
and the net reconstructed data volume will be				129.71	Tbytes

and the estimated amount of CPU needed using typical current cores is:

	Lumi in multiples of raw data cnt	#Raw Events	CPU (s)/event	CPU seconds
Recon	2	3.63E+10	0.1	7.26E+09
MC	0.1	3.63E+09	5	1.81E+09

Model for data distribution/reconstruction has been rapidly evolving

1 month ago:

- Assumed raw data would all be exported and reconstructed at SLAC/UNH
- Then informed that JLAB would not be able to export that quantity of data

Three weeks ago:

- Reconstruct at JLAB and export recon'd data only
- No go … same amount of data flow

Now:

- Reconstruct at JLAB
  - Still waiting for response on available resources
- For analyses Simulate an L3 trigger at JLAB cutting data rate to 1/3
- Export to SLAC and UNH or if this is too much then SLAC and ultimately from SLAC on to UNH
- For detector studies, small samples of full stream will be exported to SLAC and some studies will simply be run at JLAB

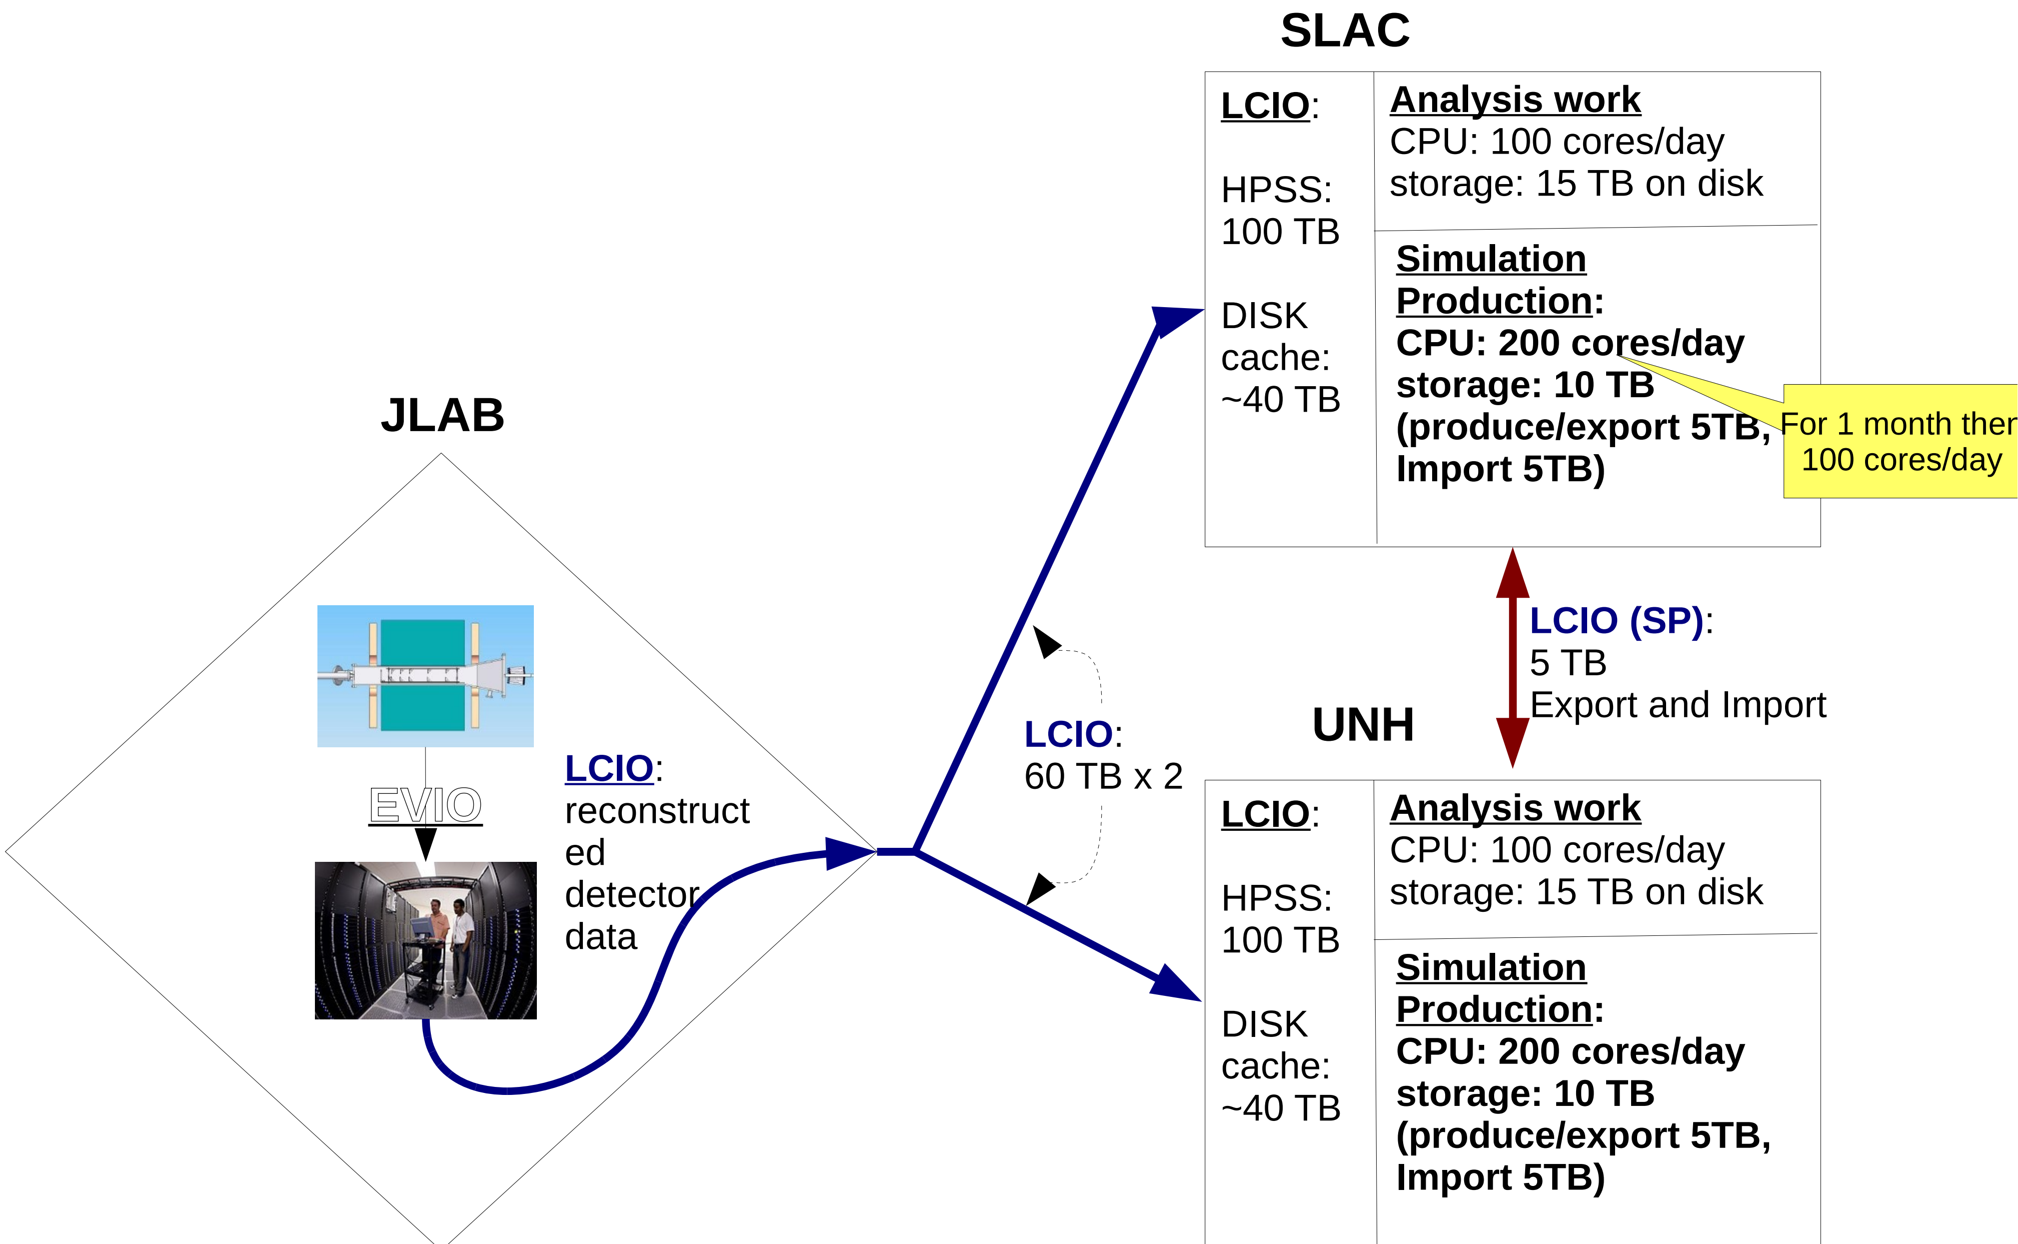

https://wiki.jlab.org/cc/external/wiki/index.php/Scientific_Computing_Systems

Current HPS Model

Difference from other models:

- Primary stream of data to SLAC and UNH just for physics studies; an L3 simulator is used to reduce data rate to ~1/3 of input rate. Small samples of unfiltered data will be provided for detector performance studies.
- More of the simulation load will be carried by SLAC w.r.t. UNH

SLAC

LCIO:

HPSS: 40 TB

DISK cache: ~20 TB

Analysis work
CPU: 100 cores/day
storage: 15 TB on disk

Simulation Production:
CPU: 300 cores/day
storage: 10 TB
(produce/export 7.5TB,
Import 2.5TB)

For 1 month then 100 cores/day

LCIO (SP):
SLAC Export 7.5TB
and SLAC Import 2.5TB

JLAB

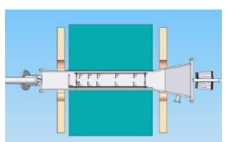

EVIO

LCIO:
60TB x 2
reconstructed
detector data

Lite L3:
export only
events of
physics interest

Small samples
of the full
stream will
also be
exported

LCIO:
20 TB x 2

UNH

LCIO:

DISK cache: ~20 TB

Analysis work
CPU: 100 cores/day
storage: 15 TB on disk

Simulation Production:
CPU: 100 cores/day
storage: 10 TB
(produce/export 2.5TB,
Import 7.5TB)

https://wiki.jlab.org/cc/external/wiki/index.php/Scientific_Computing_Systems

Data Analysis Cluster https://wiki.jlab.org/cc/external/wiki/index.php/Scientific_Computing_Systems

JLab Auger Software

JLab's Auger software manages the interface to the PBS/Maui batch farm.

Data Analysis Cluster

The batch farm is a 64-bit CentOS 5.3 environment.

The alias 'ifarml64' (interactive 64-bit linux farm) is the primary user front-end to the batch computing farm. The farm is accessible via Auger's jsub command.

```
* ifarm1101: 32-core AMD Opteron, 64GB RAM, 1TB local /scratch
* ifarm1102: 32-core AMD Opteron, 64GB RAM, 1TB local /scratch
Multithreaded code is encouraged!
Compute nodes:
farml2-10:                  8 Intel 8-core hyperthreaded, 24GB RAM, 500GB disk
farm09001-024              24 Intel 8-core hyperthreaded, 24GB RAM, 500GB disk
farm10001-024              24 Intel 8-core hyperthreaded, 24GB RAM, 500GB disk
farm11001-018              18 Intel 8-core hyperthreaded, 24GB RAM, 1000GB disk
farm11019-022               4 ADM Opteron 32-core, 64GB RAM, 1000GB disk
Work Disks
```

Work Disk (/work) servers provide ~100TB work space on SunFire X4500 series Thumper/Thor fileservers, for use in the data analysis process. Work areas are not backed up; some are allocated to a group and under its control while other areas are auto-managed to keep a pool of space available during data processing.

The Mass Storage System

JLab JASMine Software

Jasmine : controls media access and transfer of data between tape (/mss) and the stage (/stage) pools. It also manages the cache servers (/cache), which contains online copies of files from offline tape storage.

Cache Disks

Online disk storage in /cache consists of a mix of Linux and Sun X4500 (Thumper & Thor) fileservers, and by definition hold copied of data already stored on tape in the mass storage system. Some areas of cache are assgined to groups as DST (Data Summary Tape) storage, while others are automanaged.

The allocation and layout of the secondary mass storage can be found on the Cache Disk information pages.

IBM TS3500 Tape Library

An IBM TS3500 tape library, with current configuration (Fall 2011) of 6800 slots in 12 frames, with 8 LTO-4 120 MB/sec drives and 800GB cartridges, and 4 LTO-5 140 MB/sec drives with 1500GB cartridges. The total capacity is ~7PB.

HPS computing hardware needs from SLAC

LCIO: Detector Data HPSS: 40 TB DISK cache: ~20 TB	**Analysis work** CPU: 100 cores/day storage: 15 TB on disk **Simulation Production:** •**CPU:** • **300 cores/day** •**storage: 10TB** •**I/O:** • **produce/export 7.5TB,** • **Import 2.5TB**

For 1 month then 100 cores/day

LCIO (MC):
SLAC Export 7.5TB
and SLAC Import 2.5TB

HPS Computing FTE needs

	Functions	FTEs needed From now till one month after data taking	FTEs needed through the rest of the year
Norman Graf	G4, LCSIM, Java	***	**
Jeremy McCormick	LCSIM, Java	***	**
Tony Johnson ++ (chee, ...)	Tools (conditions database, forums, CVS)	*	*
Ryan Herbst	SVT DAQ	***	*

Concerns

- Raw data format is just now being settled/debated (EVIO3 vs. EVIO4, just getting our first look at the FADC output, how it is to be package in EVIO?)
- Java expertise (beyond those veterans of the LCD group, HPS members are new to Java)
- I/O – transfer capacity from JLAB to SLAC
- Potential storage costs (disk/tapes)
- Access to computational resources (shares for HPS, etc...)
- Remote management/coordination of production

Commonalities with other projects

- HPS-LCSIM is an outgrowth of the LCD lcsim effort
- Condition database is be adapted from what was done for EXO and FERMI
- Starting to use group manager software
- Starting to use SLAC Forum
- G4